Structured Sparse Methods for Active Ocean Observation Systems with Communication Constraints

Urbashi Mitra, Sunav Choudhary, Franz Hover, Robert Hummel, Naveen Kumar, Shrikanth Naryanan, Milica Stojanovic, and Gaurav Sukhatme

Urbashi Mitra, Sunav Choudhary, Naveen Kumar, Shrikanth Naryanan, and Gaurav Sukhatme are with the University of Southern California.

Franz Hover and Robert Hummel are with the Massachusetts Institute of Technology.

Milica Stojanovic is with Northeastern University.

ABSTRACT

Actuated sensor networks enabled by underwater acoustic communications can be efficiently used to sense over large marine expanses that are typically challenged by a paucity of resources (energy, communication bandwidth, number of sensor nodes). Many marine phenomena of interest admit sparse representations, which, coupled with actuation and cooperation, can compensate for being data starved. Herein, new methods of field reconstruction, target tracking, and exploration-exploitation are provided, which adopt sparse approximation, compressed sensing, and matrix completion algorithms. The needed underlying structure (sparsity/low-rank) is quite general. The unique constraints posed by underwater acoustic communications and vehicle kinematics are explicitly considered. Results show that solutions can be practically implemented, even over large ocean spaces.

INTRODUCTION

Future advances in ocean monitoring, offshore industry, and basic marine sciences will rely heavily on our ability to jointly consider communication, actuation, and sensing in a unified system that includes remote instruments, underwater vehicles, human operators, and sensors of all types. These tasks will require methods to detect and track large-scale ocean phenomena such as algal blooms, oil spills, ocean currents, and hydrothermal vents, as well as man-made signals such as those emanating from an airplane's black box. We envision a scenario as depicted in Fig. 1, where multiple autonomous underwater vehicles (AUVs) interact and coordinate via acoustic communications with a network of sensors to detect and track a phenomenon of interest. The underlying network architecture includes both static, communication-enabled sensor nodes, as well as actuated nodes in the form of AUVs. Thus, our system needs to control and move some of the nodes to achieve its sensing and communication goals. Moreover, the choices made regarding communication, control, and sensing are interdependent [1].

These goals require active consideration of the underlying problem of exploration and exploitation. A critical feature of this problem is sparsity: our sensors will be sparsely deployed due to cost and complexity considerations; our fleet of AUVs will be modest in number. As a result, our observations of a vast ocean can only be a scant sampling of the phenomena of interest. Satellite remote sensing and traditional ship operations in a large-scale underwater observation system are similarly limited, and many ocean measurements are made point-wise or confined to a local area. Furthermore, we often need to first find (explore for) a feature and then track (exploit) it, as illustrated in Fig. 2.

While there has been strong attention paid to exploiting the sparse nature of underwater acoustic channels (e.g., [2]), our problems of interest are very different. We wish to examine how to task sensors, static and/or mobile, toward achieving an underwater mission goal that will be dependent on the sensor network's understanding of its environment. Prior work related to field estimation and target tracking suggests that sparse representations exist for many features of interest, and therefore, we may not be irreparably data-starved. At the same time, a mission has to be accomplished by coordinated sensing and action, typically enabled by acoustic communication between energy-constrained and kinematically constrained agents. In this article, we review how sparsity can be leveraged in the design of exploration-exploitation methods tailored to underwater sensing systems. Our particular approach does not require detailed prior information about the fields or targets of interest.

In particular, we tackle three related problems: field reconstruction [1, 10, 11, 12], target tracking [5, 13], and methods for exploration-exploitation [14, 15], which build on the insights derived from the previous approaches. Field reconstruction is an important issue for many scientific, industrial, and tactical applications. Examples include geographical, chemical, and temperature map building. In some cases, reconstruction is the end goal (e.g., climate recording), while in other cases it represents an intermediate step to support the mission at hand, for example, one in which mobile assets need to learn about the environment in which they operate and adapt to it. Sensor measurements of the field are typically correlated in some domain, and this correlation forms the basis of compressive sensing theory: when the field to be reconstructed is sparse (with respect to some basis), it suffices to collect data only at some of the points to reconstruct the full view of the field.

In recent years, there has been a surge of new results in the area of sparse field reconstruction using sensor networks. Although harsh environments, energy management, and limited communications have been considered, most of these methods have been optimized for terrestrial conditions. Underwater systems pose new challenges and constraints. For example, sparse methods have been developed predominantly ignoring the cost of sampling, while in oceanographic field applications, constraints such as limited energy and maneuverability can dominate, and the number of samples may be quite low. The interplay between such resource constraints and the fundamental sparse algorithms is a general problem we refer to as kinematically constrained sparse approximation (KCSA) [10, 11].

The extension from field reconstruction to target tracking is a natural one, and we provide methods that explicitly consider the dynamics of the targets and the nature of the underwater acoustic communication channel. An example is random access compressed sensing (RACS) [12, 13], which notably does not require synchronization. Target detection is a related application that addresses whether a specific phenomenon of interest has occurred. The target usually cannot be observed directly, however, and the task then is to infer the presence or location of the target from indirect observations, with an eye to the cost of acquiring each sample, thus resulting in the exploration-exploitation [14, 15] optimization trade-off.

We summarize key challenges of underwater sensing and communication networks that are enabled by actuation, including the unique features of the acoustic channel and the constraints of underwater vehicles. We lay out the key ideas in compressive sensing, sparse approximation, and matrix completion. We examine how to reconstruct a large field when limited to a few mobile vehicles operating with constraints over a large space. We change the perspective by considering field reconstruction when using static underwater sensors. Here, the impact of using the underwater acoustic medium to transmit sensor data to a fusion center introduces unique complexities; and the work is extended to the tracking of mobile sources. Finally, the insights of the prior two research areas are employed to derive novel methods based on matrix completion to solve a realistic exploration-exploitation problem.

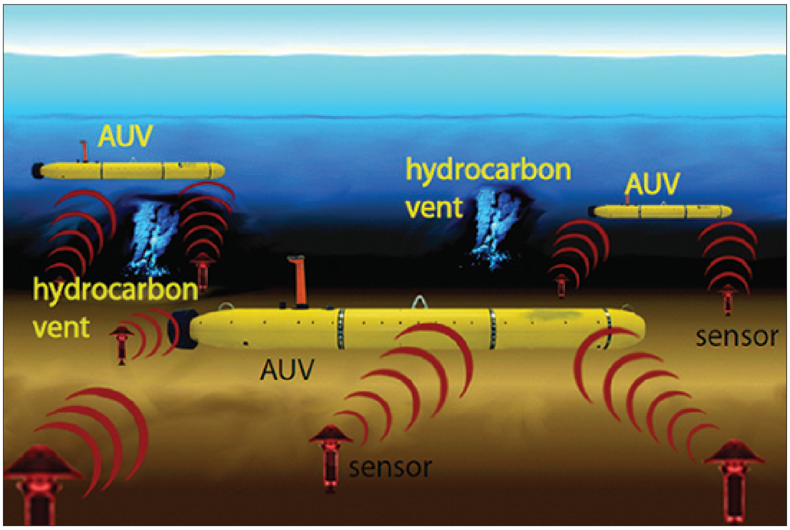

Figure 1. The envisioned network architecture for future ocean observing systems, comprising both fixed and moving nodes capable of advanced sensing and wireless communications.

The Challenges of Underwater Actuated Networks

The boundaries between communications, networking, navigation, and sensing are blurred in underwater actuated sensor networks, where multiple levels of dynamics exist: moving (possibly actuated) nodes, time-varying communication channels, and evolving phenomena to be detected, identified, and tracked.

Underwater Acoustic Communication Channels

Challenges in the design of underwater acoustic communication systems include a severely limited, range-dependent bandwidth, extensive time-varying multipath propagation, and long propagation delays caused by the low speed of sound underwater (1500 m/s) [3]. Absorption leads to exponential path losses with respect to frequency. Propagation delays result in significant delay spreads (on the order of hundreds of symbols). There are very short channel coherence times that are exacerbated in mobile environments and orders of magnitude shorter than in RF systems. In terms of the channel, we focus on the sparse nature of multipath propagation.

Underwater Sensors and Vehicles

Ocean vehicle systems face obstacles not shared in the terrestrial and aerial domains. Most of the ocean remains uncharted, with detailed information about the sea floor available only at specific and scientifically important sites, or areas of specific interest such as oil and gas fields. Underwater systems are expensive to build and instrument, and ship time also introduces a huge cost that encourages remote operations. Unlike land-based vehicles, underwater vehicles are exposed to extreme pressure, corrosion, and fouling, and strong variable wave, current, and wind disturbances. Such disturbances make operating a vehicle difficult, if not impossible. At the

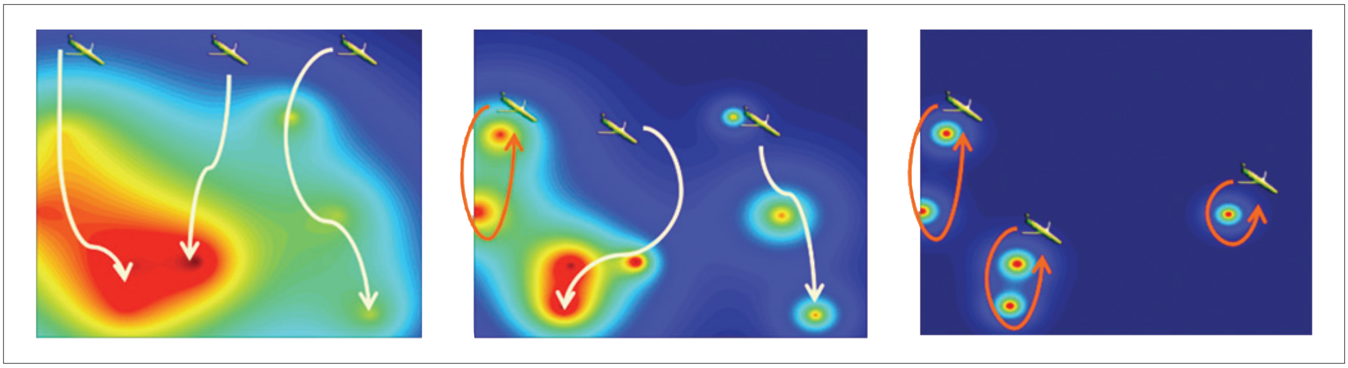

Figure 2. Exploration-exploitation at sea: autonomous agents first explore in a coordinated way to identify phenomena of interest (left); as the field evolves, the AUV network differentiates to allow individual agents to follow particular targets (exploitation).

same time, it should be noted that currents can sometimes be exploited opportunistically, since energy consumption is a major concern from the point of view of propulsion, acoustic communications, and hotel load. Aerial vehicles face similar energy issues, but usually have two key advantages over underwater systems: low-power communication over an RF channel with other vehicles or a base station, and the ability to measure their own location with precision using GPS. Underwater positioning methods are often acoustic and subject to the challenges noted above for acoustic communications. Today, ocean vehicles typically break the surface for precision localization via GPS — a hazardous and time-consuming operation.

As one specific example of these challenges, we carried out an experiment with an autonomous surface vehicle (ASV) station-keeping via commands sent through an acoustic modem. The effects of wind in such a setting are to both push the boat off the reference point and degrade the acoustic channel because of surface chop. Indeed, we observed a full order of magnitude increase in positioning error as the wind speed increased from 1 to 8 m/s [4]. Another experiment coordinating vehicles through acoustic communication is detailed in [5]. Overall, the costs of actuation, communication, and sensing can be comparable for marine systems [3], and this parity strongly affects how underwater acoustic networks are to be optimized in support of a mission goal [1].

Sparse Approximation and Matrix Completion Review

Here we review key elements of sparse approximation and matrix completion needed for understanding the methods presented in this article.

Any family of signals with a linear algebraic representation requiring many more parameters than the actual number of unknowns in a given instance is called a low-dimensional signal. Two low-dimensional signals that have been the subject of active research are sparse vectors and low-rank matrices.

Considering sparse vectors first, let $\Psi \in \mathbb{R}^{n \times n}$ be an orthonormal basis so that any signal $f \in \mathbb{R}^n$ can be written as $f = \Psi x$ for a coefficient vector $x \in \mathbb{R}^n$. x is called *s-sparse* if it has at most s nonzero elements, and if a sparse x suffices, the signal f is said to admit a sparse representation in Ψ. Intuitively, an s-sparse signal has only s unknowns, many fewer than the full representation [6]. f is called compressible if it is well approximated — but not constructed exactly — by a sparse x. Many natural phenomena are considered to be compressible in domains such as frequency, wavelets, or total variation. For low-rank matrices, a key classical result is that any $m \times n$ matrix X admits the singular value decomposition $X = U\Sigma V$, and that the number of nonzero elements in Σ denotes the rank of X. The Eckart-Young-Mirsky result says further that reduced-rank approximations to Σ, obtained simply by keeping only the r largest elements in Σ, are optimal in the Frobenius norm [7]. Low-rank matrices arise frequently in second-order datasets (Netflix) and data generated from bilinear observations [8].

Incoherence is a critical requirement for efficiently sampling low-dimensional models, and measures the dissimilarity between the signal basis and the measurement basis. Information measured in one basis will "spread out" when measured in a second (incoherent) basis, enabling efficient reconstruction even if few measurements are made. In point-wise compressed sensing, the measurement $y = R\Phi f$ involves a "fat" matrix $R\Phi$, where R is a fat projection (a row-reduced identity matrix), and Φ an orthonormal measurement basis. Coherence of the measurement and signal bases is defined as

$$\mu(\Phi,\Psi)=\sqrt{n}\max_{1\le j,k\le n}\left|<\phi_j,\psi_k>\right|,$$

where ϕ_j and ψ_k denote the j^{th} and k^{th} atoms of Φ and Ψ, respectively, and $<\cdot,\cdot>$ denotes an inner product. For a 2D field represented in the discrete cosine basis, and sampled with delta functions, $\mu(\Phi,\Psi) = 2$, a highly incoherent basis pairing is indicated. The reconstruction task can be formulated as a convex optimization [6]

$$\hat{x} = \operatorname{argmin} ||z||_1 \text{ subject to } y = R\Phi\Psi z,$$

where $||\cdot||_1$ is the one-norm. Pulling together sparsity and incoherence, a key theorem of compressed sensing states that if f is s-sparse in Ψ, then $\hat{x} = x$ with high probability if the number of measurements (i.e., the number of rows in R) exceeds $C\mu^2(\Phi,\Psi)s \log n$ for some positive con-

Figure 3. Kinematically-constrained sparse approximation in action at sea. The left image shows two "random TSP" paths generated and executed by an autonomous surface vehicle (center); the right image shows a depth-field reconstruction based on sparse samples and the discrete-cosine transform.

stant C. Coherence, and thus the associated sampling rate, determine the feasibility of the reconstruction.

The measurement ideology in low-rank matrix completion bears a strong resemblance to that of compressed sensing. Here $y = \Omega(X)$, where Ω denotes a small set of indices, and hence the measurement basis is the set of delta functions over the vector space of $m \times n$ matrices. Incoherence is defined as the sum of the inner products, squared, between delta functions and the left and right singular vectors [7]. The reconstruction is formulated as the convex problem

$$\hat{X} = \text{argmin } ||Z||_* \text{ subject to } y = \Omega(Z),$$

where $||\cdot||_*$ is the sum of singular values. Akin to compressed sensing, if Ω is sampled uniformly at random, $\hat{X} = X$ with high probability when the number of measurements exceeds $Cv(X)rn \log^2 n$, where r is the rank of X, and $n > m$.

Turning to the selection of points via R, random sampling gives strong stability guarantees for both compressed sensing and low-rank matrix completion in the presence of measurement noise as well as model mismatch. Mismatch occurs when the unknown signal is compressible instead of sparse, or is approximately low-rank instead of exactly low-rank. In compressed sensing, we appeal to the restricted isometry property (RIP) of the measurement matrix of order s, by which the measurement process approximately preserves pairwise distances between s-sparse vectors. In low-rank matrix completion, a weaker condition called *restricted strong convexity* is desirable, and similarly sufficient to show stability, given $r \ll n$ [9].

Random sampling may be infeasible or unsuitable, and it is desired in such cases to find new projection matrices that:
- Result in efficient reconstruction
- Can be implemented efficiently by the sampling system (e.g., mobile robots in a large domain)

In general, it is not easy to show that structured projections satisfy the RIP, although there have been recent efforts directed toward the design of deterministic projection matrices. These approaches control the structure of the sensing matrix $R\Phi\Psi$; however, in our case we can only control the measurement matrix $R\Phi$. To the best of our knowledge, there are no techniques for structured random construction of the measurement matrix, especially given the motion constraints of robots.

Field Reconstruction with Mobile Sensors via KCSA

Since kinematic constraints are one of our main divergences from standard sparse approximation, a natural step is the construction of a cost related to resources. We can identify three key aspects of such a cost:
- The Euclidean length of a transit between two sample sites
- A heading change (maneuvering) required to move between two successive sites
- The resources required for an individual sample itself

Working with a large sea-surface temperature dataset and the discrete cosine transform basis, we optimized statistics of randomly drawn transit lengths and maneuvers to minimize the sparse approximation reconstruction error in a simple random walk scheme [10]. This optimized random walk, surprisingly, was outperformed in sparse approximation error by an even simpler procedure of randomly choosing a set of sites and connecting them with a nearest-neighbor traveling salesperson problem (TSP) solution. Figure 3 depicts an experimental test with this method. More generally, reconstruction errors can be bounded by ensuring that the RIP in the sampling basis is achieved, suggesting a second class of lightweight KCSA methods. For a given dictionary and a set of sample points, restricted isometry is trivial to compute incrementally, so a vehicle can be instructed to move to randomly selected, as yet unvisited points based simply on what the next point will do to the RIP. This is a purely greedy approach that typically offers an improvement of 10–20 percent relative to points selected without considering the RIP [11].

Maintaining restricted isometry and minimizing path length could be tackled as well in a more formal optimization setting. With the spike basis, however, the projection (sampling) design is a binary programming problem. Furthermore,

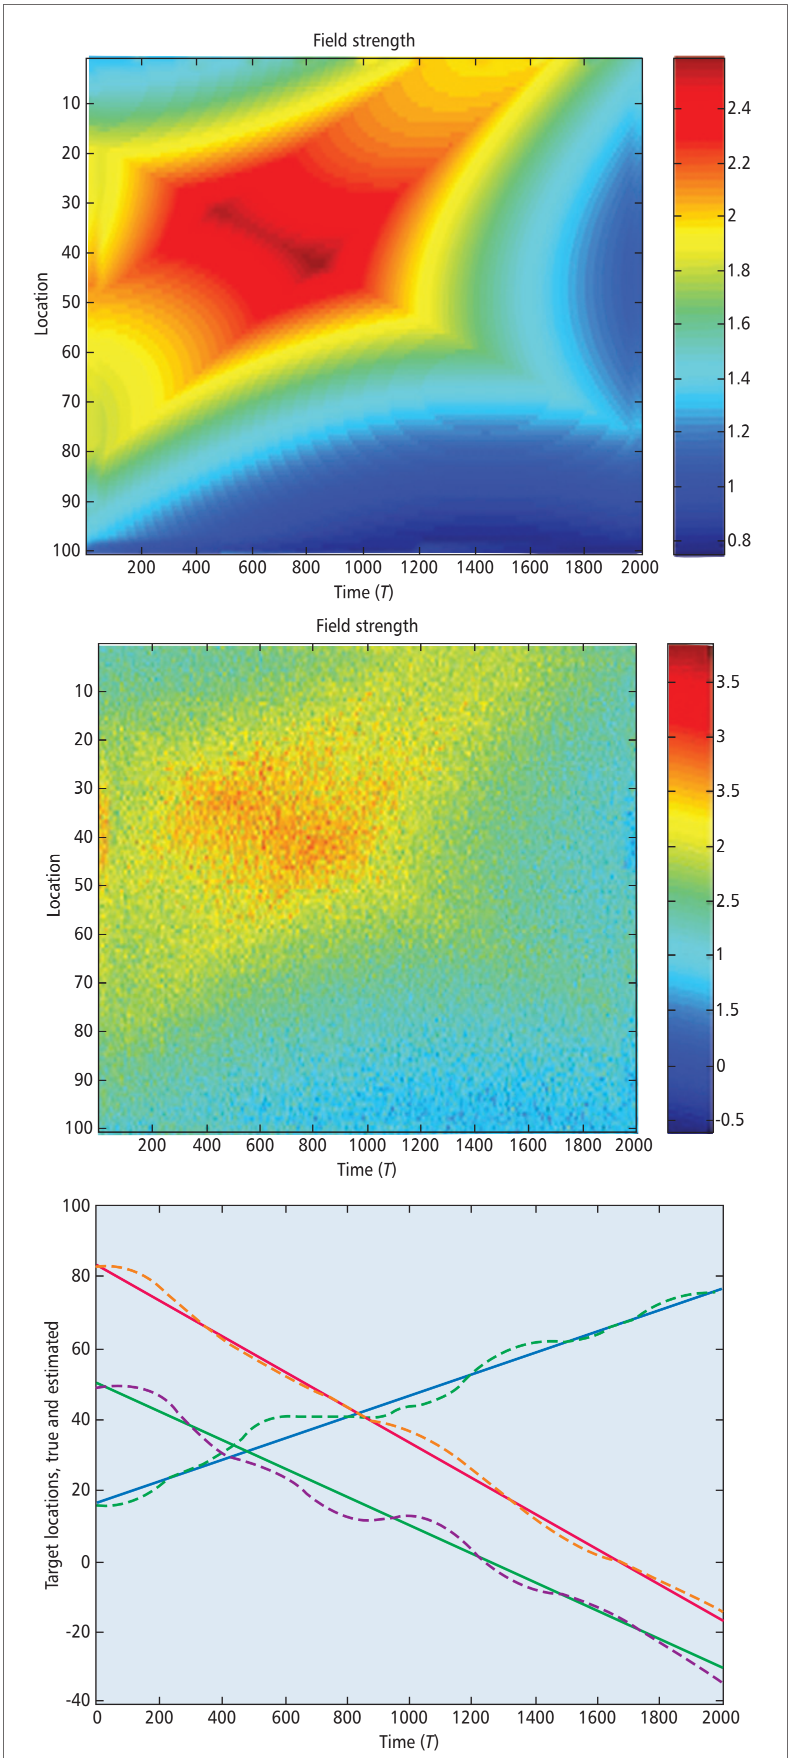

Figure 4. Target tracking without field reconstruction for three mobile sources. Instaneous observations of the noise-free (top) and noisy (center) fields; (bottom) true and estimated paths.

we have found that when the number of samples is small, there is virtually no advantage in the designed plan over random points; with a larger number of samples, the designed points have a coherence typically 20–30 percent lower than the random points. The formal sampling problem can also be merged with the TSP, but complexity scaling is very poor since they are not well coupled. An alternative is to generate a pool of optimal sample sets without attention to the path, and then apply strong TSP solvers to members of the pool. Applying such a process, we find that a pool typically identifies a Pareto-like frontier of TSP cost vs. coverage (e.g., as characterized by the star discrepancy), and thus a trade-off space between sampling efficacy for the purpose of reconstruction and the path length. This multi-way trade-off is developed further later.

Field Reconstruction in an Underwater Sensor Network

As noted earlier, in underwater acoustic systems, key physical constraints include limited bandwidth and the high rate of packet loss. Moreover, synchronization is difficult, often mandating a design that does not rely on accurate timing. To address these issues, we draw on the sparse (compressible) nature of the sensing field, through a design that capitalizes on random channel access and compressive sensing (RACS), using multiple sensors interacting with a fusion center. In this approach, a sensor takes local field measurements at random instants in time [12]. Each measurement is encoded into a data packet, which is immediately transmitted to the fusion center. The fusion center collects all the packets over a certain interval of time, and then attempts to reconstruct the field.

The key (and perhaps counterintuive) idea behind combining random access and compressive sensing is that the fusion center does not care which sensors provide the data packets as long as the sensors are selected uniformly at random, and there are sufficient packets. With random access, packet collisions are inevitable. However, they occur randomly and thus do not alter the way in which the FC perceives packet arrivals. Consequently, the fusion center can simply discard the packets that have collided (as well as those that were received in error due to noise, fading or interference), and perform field reconstruction using the remaining useful packets. Because of the random nature of this architecture, a probabilistic approach to system design becomes necessary. Our approach relies on the notion of sufficient sensing probability, the probability with which full field reconstruction is achieved in a certain interval of time. Setting this probability to a desired target value, system optimization under a relevant criterion (e.g., minimum energy per bit or minimum reconstruction time) yields the necessary parameters: the packet generation rate (per-node sensing rate), required bandwidth, and power. Many fields of interest are sparse with respect to fairly general dictionaries; thus, the reconstruction method is relatively non-parametric.

The concept of random access compressive sensing can also be extended to the tracking of vehicles each of which emanates a particular signature [13]. Based on the knowledge of signatures, a reconstruction algorithm is designed that side-steps the field reconstruction, and instead focuses directly on object tracking. Figure 4 provides an example of the multi-vehicle tracking capabilities of RACS. Clearly, there is correlation in the map data (i.e., the field is sparse). This particular spatio-temporal field is created by three moving objects, each having an exponentially decaying signature; the level of sparsity is three. Shown on the right are the objects' true trajectories (solid) and estimates obtained by the tracking algorithm (dashed), applied to samples of the noisy field (–20 dB signal-to-noise ratio, SNR, 30 percent packet loss) obtained from 10 sensors distributed uniformly along a track. As with field reconstruction, effective target tracking based on RACS does not require complex models of vehicle behavior (vs. field behavior). As such, our methods are highly robust.

Due to the time-varying nature of the field, one needs to optimize the sensing duration — the key is to acquire a sufficient number of samples for reliable reconstruction, while ensuring that the sensing duration is short enough so that the observed field does not change significantly within a single sensing period. From the perspective of target detection using the acquired image, there is another notion of sensing rate optimality corresponding to the trade-off between detection accuracy and energy efficiency. One must also consider the rate at which the field is changing due to the speed of the vehicles to be tracked. To address these issues, a feedback scheme can be employed, resulting in adaptive sensing. The feedback signal is generated by performing target detection on the acquired field, which is then compared to a model of the field, resulting in a model-based error. Figure 5 depicts the evolution of the sensing rate (T) based on such an algorithm. If there is a large model-based error, the sensing rate is increased (state A). To compensate for too high a sensing rate, a motion-dependent error measure is also computed (state D). If the vehicles move slowly, the fusion center can respond by decreasing the sensing rate to conserve energy. Thresholds on the two error measures are optimized to maximize performance while constraining the sensing energy. States B and C correspond to conflicts in the decisions resulting from the two different sensing metrics.

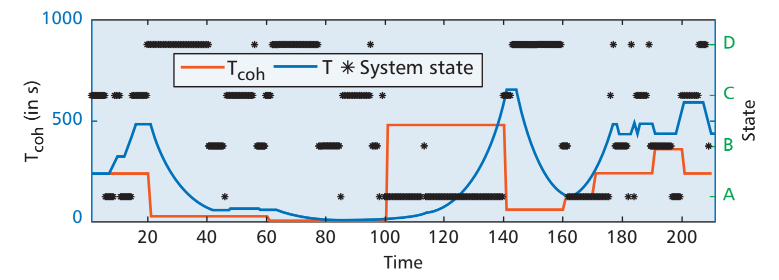

Figure 5. Controlled sensing rate T using the proposed method; states A and D correspond to undersensing and oversensing, respectively; states B and C correspond to the cases when the model error and sensing error metrics lead to conflicting decisions. (T_{coh} is the coherence time of the time-varying field).

Tackling the Integrated Exploration-Exploitation Problem

Our prior work on field reconstruction shows that kinematic and communication constraints can readily be employed to develop practical reconstruction methods exploiting new results in sparse approximation. Furthermore, a large family of fields admits the needed sparsity. We can extend these notions to the exploration-exploitation problem. As a concrete example of integrated search with sparse methods, we consider an underwater target that emits a field of observable intensity, as in Fig. 6a; close to the target, the intensity is high, but the intensity decays as we move away from the target location. The aggregate field shown here is a side-scan sonar image, and the target signature could be due to an acoustic signal from a flight data recorder lost in some sector of the ocean. A deployed AUV is restricted to collecting a few samples along some navigation path (Fig. 6b), with the hope that these samples contain enough information to localize the target to a smaller region of interest (Fig. 6c). The intensity field in Fig. 6a is approximately separable, however, and therefore admits a low-rank matrix structure with respect to the image matrix. This property allows the reduction of the number of samples required for target localization [14], subsequently leading to reduced cost of navigation [15]. On the other hand, a decaying low-rank target field is also a very generic structural assumption not specific to this particular dataset. Thus, a rich set of target signatures meet the two key assumptions; additionally, we can extend our results to allow for multiple targets, different adaptive and non-adaptive sampling approaches, and even distributed implementations using multiple AUVs with communication costs. With our approach, we do not need to know the target signature *a priori*.

We developed a new methodology for this problem based on low rank matrix completion and the high-dimensional properties of the traveling salesperson problem, and one appealing outcome is a characterization of a multi-way trade-off between sampling, navigation, computation, and localization [15]. Our algorithm takes a hierarchical exploration-exploitation approach that combines the strengths of low-rank matrix completion and binary search, viz. noise robustness and computational speedup. In each stage, based on the current field of view, a set of random locations is selected; the AUV determines an optimized path covering these locations, at which sample observations are collected. The field is then reconstructed using low-rank matrix completion [7], computing the best rank one approximation. After post-processing, the target is localized, a smaller field of view is determined, and the steps are then repeated. We assume that the algorithm terminates after k

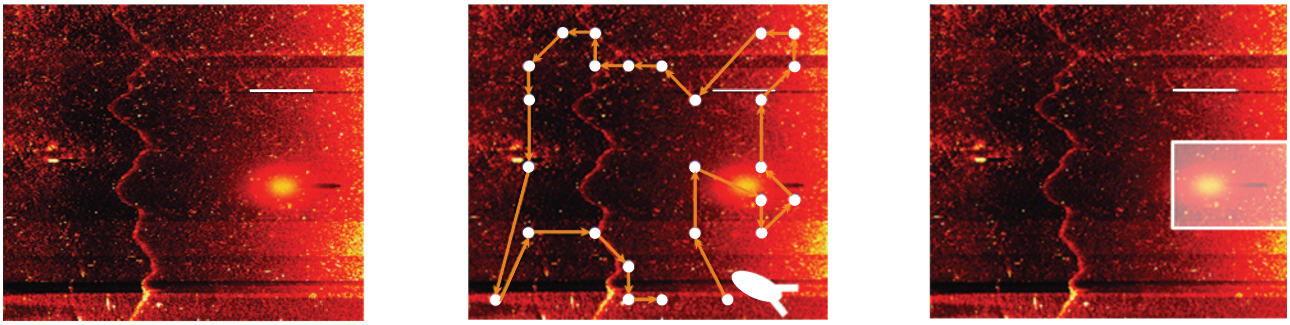

Figure 6. (left) Original intensity map of an underwater target (sonar data); (center)initial sampling path of the AUV; (right) the localized region of interest after several stages of the matrix completion exploration-exploitation algorithm.

iterations, define the total search space of size $n \times n$, and the field of view to be of size $m \times m$ representing a uniformly decimated sub-matrix of the total search space. With these definitions, theoretical results on low rank matrix completion and Euclidean TSP can be used to show that our approach takes $O(km \log^2 m)$ samples and $O(km^{3.5})$ computational time and achieves $O(k\sqrt{m} \log m)$ cost of navigation with high probability, where the number of iterations increases with noise level and satisfies $\log_m n \leq k \leq n/m$.

A design parameter is the field of view. Intuitively, the size of the field of view is chosen inversely proportional to the spread in the target field since a larger spread enables collecting sample intensities above the noise floor farther from the target location, thus reducing sampling requirements. Decaying separable fields admit good rank-one approximation via singular vectors that are also decaying. The decaying property of the singular vectors can be used for further processing via unimodal function fitting (a form of isotonic regression), which results in greater noise reduction and improved target localization. Numerical results show that there is an optimum spreading factor leading to the best detection performance. One can additionally characterize the best balance between existence of a significant gradient in the signature (after noise corruption) and the target signature occupying a larger footprint; that is, fields with large support of the gradient show better performance for a fixed sample complexity. We can also provide a theoretical analysis that provides the trade-off between our ability to localize the target within a stage and our ability to reconstruct the current windowed field. Finally, the overall multi-way trade-off can be characterized by including navigation and kinematic costs (or constraints) as well as the incorporation of the randomness induced by packet dropping in the underwater acoustic communication channel. Thus, the features of our field reconstruction, tracking, and exploration-exploitation work can be seamlessly combined.

Conclusions

We have argued that mobile underwater sensor networks enabled with control and acoustic communication will play a major role in future ocean observing infrastructure. Such autonomous systems need novel mechanisms for handling underwater communication constraints, and need to integrate sensing and classification in providing solutions for key exploration-exploitation trade-offs. The unavoidable burden of limited resources in these operations (e.g., physical energy, time, communications, and accurate synchronization) are in a sense balanced by the promise of compressed sensing and low-rank approximation theory over huge spatial domains, coupled with standard tools including solvers for the traveling salesperson problem. Thus, despite the paucity of resources and limited number of deployed sensors, whether they be static or mobile, the underlying sparsity (or low rank property) of the processes of interest enables us to overcome the challenge of being data-starved. Our structures of interest are general and do not require complex a priori model information for our fields or targets. We have specifically considered constraints inherent to our underwater system: the kinematic constraints of vehicles as well as the unique characteristics of the underwater acoustic channel. These constraints have led to novel problem formulations and attendant solutions. In fact, some of our most recent outcomes can bypass weaknesses of more classical approaches to solving the exploitation-exploration problem that are not flexible to incorporating multiple resource constraints such as navigation, communication, and sensing while allowing for theoretical analysis. The successful implementation of these methods at sea will have an impact on many applications and industries, including environmental monitoring, aquatic ecosystems, ocean accident remediation, surveillance for defense applications, homeland security, the oil and gas industry, aquaculture, geological and oceanographic science, and marine biology.

References

[1] G. Hollinger *et al.*, "Underwater Data Collection Using Robotic Sensor Networks," *IEEE JSAC*, Special Issue on Communications Challenges and Dynamics for Unmanned Autonomous Vehicles, vol. 30, no. 5, June 2012, pp. 899–911.

[2] C. Berger *et al.*, "Sparse Channel Estimation for Multicarrier Underwater Acoustic Communication: From Subspace Methods to Compressed Sensing," *IEEE Trans. Signal Processing*, vol. 58 no. 3, Mar. 2010, pp. 1708–21.

[3] M. Stojanovic and J. Preisig, "Underwater Acoustic Communication Channels: Propagation Models and Statistical Characterization," *IEEE Commun. Mag.*, vol. 47, no. 1, Jan. 2009, pp. 84–89.

[4] F. Arrichiello *et al.*, "Effects of Underwater Communication Constraints on the Control of Marine Robot Teams," *Proc. IEEE Int'l. Conf. Robot Commun. and Coordination*, Odense, Denmark, Mar. 2009.

[5] B. Reed *et al.*, "Multi-Vehicle Dynamic Pursuit Using Underwater Acoustics," *Proc. Int'l. Symp. Robotics Research*, Dec. 2013, Singapore.

[6] E. Candes and T. Tau, "Decoding by Linear Programming," *IEEE Trans. Info. Theory*, vol. 51, no. 125, Dec. 2005, pp. 4203–21.

[7] D. Gross, "Recovering Low-Rank Matrices From Few Coefficients in Any Basis," *IEEE Trans. Info. Theory*, vol. 57, no. 3, Mar. 2011, pp.1548–66.

[8] S. Choudhary and U. Mitra, "On Identifiability in Bilinear Inverse Problems," *Proc. IEEE Int'l. Conf. Acoustics, Speech and Signal Processing*, May 2013, Vancouver, Canada, pp. 4325–29.

[9] S. Negahban and M. J. Wainwright, "Restricted Strong Convexity and Weighted Matrix Completion: Optimal Bounds with Noise," *J. Machine Learning Research*, vol. 13, no. 1, Jan. 2012, pp. 1665–97.

[10] R. Hummel *et al.*, "Mission Design for Compressive Sensing with Mobile Robots," *Proc. Int'l. Conf. Robotics and Automation*, May 2011, Shanghai China, pp. 2362–67.

[11] F. Hover *et al.*, "One-step-Ahead Kinematic Compressive Sensing," *Proc. IEEE GLOBECOM Wi-UAV Wksp.*, Dec. 2011, Houston, TX, pp. 1314–19.

[12] F. Fazel, M. Fazel, and M. Stojanovic, "Random Access Compressed Sensing over Fading and Noisy Communication Channels," *IEEE Trans. Wireless Commun.*, vol. 12, no. 5, May 2013, pp. 2114–25.

[13] K. Kerse, F. Fazel, and M. Stojanovic, "Target Localization and Tracking in a Random Access Sensor Network," invited paper, *47th Asilomar Conf. Signals, Systems and Computers*, Nov. 2013, Pacific Grove, CA, pp. 103–07.

[14] S. Choudhary *et al.*, "Active Target Detection with Mobile Agents," IEEE Int'l. Conf. Acoustics, Speech and Signal Processing, May 2014, Florence, Italy, pp. 4185–89.

[15] S. Choudhary *et al.*, "Active Target Detection with Navigation Costs: A Randomized Benchmark," invited paper, *52nd Annual Allerton Conference on Commun., Control, and Computing*,Oct. 2014, Monticello, IL, pp. 109–15.

Biographies

URBASHI MITRA [S'88, M'98, SM'04, F'07] received B.S. and M.S. degrees from the University of California at Berkeley and her Ph.D. from Princeton University. Prior to her Ph.D. studies, she was a member of technical staff at Bellcore. After a six-year stint at the Ohio State University (OSU), she joined the Department of Electrical Engineering at the University of Southern California (USC), Los Angeles, where she is currently a professor. She is the inaugural Editor-in-Chief of *IEEE Transactions on Molecular, Biological and Multi-Scale Communications*. She is a Distinguished Lecturere for the IEEE Communications Society for 2015–2016. She was a member of the IEEE Information Theory Society's Board of Governors (2002–2007, 2012–2014) and the IEEE Signal Processing Society's Technical Committee on Signal Processing for Communications and Networks (2012–2014). She is the recipient of: a 2015 Insight Magazine STEM Diversity Award, 2012 GLOBECOM Signal Processing for Communications Symposium Best Paper Award, 2012 U.S. National Academy of Engineering Lillian Gilbreth Lectureship, USC Center for Excellence in Research Fellowship (2010–2013), the 2009 DCOSS Applications & Systems Best Paper Award, Texas Instruments Visiting Professor (Fall 2002, Rice University), 2001 Okawa Foundation Award, 2000 OSU College of Engineering Lumley Award for Research, 1997 OSU College of Engineering MacQuigg Award for Teaching, and a 1996 National Science Foundation CAREER Award. She currently serves on the following IEEE award committees: Fourier Award for Signal Processing, James H. Mulligan, Jr. Education Medal, and the Paper Prize Award. She has been/is an Associate Editor for *IEEE Transactions on Signal Processing* (2012–), *IEEE Transactions on Information Theory* (2007–2011), *IEEE Journal of Oceanic Engineering* (2006–2011), and *IEEE Transactions on Communications* (1996–2001). She has co-chaired the Technical Program Committees of the 2014 IEEE International Symposium on Information Theory in Honolulu, Hawaii, 2014 IEEE Information Theory Workshop in Hobart, Tasmania, the IEEE 2012 International Conference on Signal Processing and Communications, Bangalore India, and the IEEE Communication Theory Symposium at ICC 2003 in Anchorage, Alaska. She served as General Co-Chair of the first ACM Workshop on Underwater Networks at Mobicom 2006, Los Angeles, California. She was Tutorials Chair for IEEE ISIT 2007 in Nice, France, and Finance Chair for IEEE ICASSP 2008 in Las Vegas, Nevada. She has held visiting appointments at the Delft University of Technology, Stanford University, Rice University, and the Eurecom Institute. She served as co-Director of the Communication Sciences Institute at USC from 2004 to 2007. Her research interests are in wireless communications, communication and sensor networks, underwater acoustic communication, biological communication systems, detection and estimation, and the interface of communication, sensing, and control.

SUNAV CHOUDHARY is a Ph.D. candidate in the Ming Hsieh Department of Electrical Engineering at USC. He received his B.S. degree in 2010 from the Indian Institute of Technology, Kharagpur. He is the recipient of the following awards and honors: 2015 USC Center for Applied Mathematical Sciences Prize Winner, 2014 IEEE International Symposium on Information Theory Student Travel Grant, 2013 Ming Hsieh Institute Electrical Engineering Research Festival Best Poster Award, 2010–2014 USC Annenberg Graduate Fellowship, 2006–2010 IIT Kharagpur Goralal Syngal Memorial Scholarship, and the 20016–2010 CBSE Merit Scholarship for Professional Studies. His present research interests are in the field of sparse signal approximation and its applications to underwater acoustic communications.

FRANZ S. HOVER [M'92] received his B.S.M.E. degree from Ohio Northern University, and his M.S. and Sc.D. degrees from the Joint Program in Applied Ocean Physics and Engineering at the Woods Hole Oceanographic Institution and Massachusetts Institute of Technology (MIT), Cambridge. He was a consultant to industry and then a member of the research staff at MIT, where he worked in fluid mechanics, biomimetics, and marine robotics. His research has led to commercial development of the HAUV platform for autonomous ship hull inspection, advances in computational tools for power systems, and innovations in subsea flow control technology. He is currently an associate professor at the MIT Department of Mechanical Engineering, with research focusing on design methodologies and complex marine systems. He has authored or co-authored over 100 refereed papers.

ROBERT HUMMEL [S'11] received his B.S. and M.S. degrees in Mechanical Engineering from MIT in 2009 and 2012, respectively, with a focus on robotics, mechatronics, and autonomous sampling techniques. He experience includes designing and building autonomous vehicles, machines, and mechanisms, and various electro-mechanical systems. He has also worked with numerical simulations and optimization. Currently, he is co-founder and design engineer at ATSE in Northborough, Massachusetts, an engineering consulting firm specializing in electric motors, generators, and drives for industrial, defense, and renewable energy applications.

NAVEEN KUMAR [S'10] received his B.Tech. degree in instrumentation engineering from the Indian Institute of Technology, Kharagpur, in 2009. He is currently working toward his Ph.D. degree at the Department of Electrical Engineering, USC. His research interests include machine learning and signal/image processing for applications to speech and multimedia problems. He was awarded the Viterbi School Dean's Doctoral Fellowship at USC in 2009. He has also worked on teams that have won the Interspeech 2012 Speaker Trait Challenge and the Northern Digital Inc. Excellence Awards at the 2014 International Seminar on Speech Production (ISSP).

SHRIKANTH S. NARAYANAN [S'88, M'95, SM'02, F'09] received his M.S., Engineer, and Ph.D. degrees in electrical engineering from the University of California Los Angeles in 1990, 1992, and 1995, respectively. From 1995 to 2000 he was with AT&T Labs-Research, Florham Park, New Jersey, and AT&T Bell Labs,Murray Hill, New Jersey, first as a senior member and later as a principal member of technical staff. He is currently a professor of engineering at USC, where he directs the Signal Analysis and Interpretation Laboratory. He holds appointments as a professor of electrical engineering, computer science, linguistics, and psychology, and is the founding director of the Ming Hsieh Institute, Ming

The successful implementation of these methods at sea will have an impact on many applications and industries, including environmental monitoring, aquatic eco-systems, ocean accident remediation, surveillance for defense applications, homeland security, the oil and gas industry, aquaculture, geological and oceanographic science, and marine biology.

Hsieh Department of Electrical Engineering, USC. He has published more than 600 papers and has been granted 16 U.S. patents. His research interests include human-centered signal and information processing and systems modeling with an interdisciplinary emphasis on speech, audio, language, multimodal and biomedical problems, and applications with direct societal relevance. He is a Fellow of the Acoustical Society of America and the American Association for the Advancement of Science, and a member of Tau Beta Pi, Phi Kappa Phi, and Eta Kappa Nu. He is also an Editor of the *Computer Speech and Language Journal* and an Associate Editor of *IEEE Transactions on Affective Computing*, *APSIPA Transactions on Signal and Information Processing*, and the *Journal of the Acoustical Society of America*. He was also previously an Associate Editor of *IEEE Transactions on Speech and Audio Proseccing* (2000–2004), *IEEE Signal Processing Magazin*e (2005– 2008), and *IEEE Transactions on Multimedia* (2008–2011). He has received a number of honors, including Best Transactions Paper awards from the IEEE Signal Processing Society in 2005 (with A. Potamianos) and 2009 (with C. M. Lee), and selection as an IEEE Signal Processing Society Distinguished Lecturer for 2010–2011. Papers coauthored with his students have won awards at Interspeech 2014 Cognitive and Physical Load Challenge, Interspeech 2013 Social Signal Challenge, Interspeech 2012 Speaker Trait Challenge, Interspeech 2011 Speaker State Challenge, Interspeech 2013 and 2010, Interspeech 2009 Emotion Challenge, IEEE DCOSS 2009, IEEE MMSP 2007, IEEE MMSP 2006, ICASSP 2005, and ICSLP 2002.

MILICA STOJANOVIC [(SM'08, F'10] graduated from the University of Belgrade, Serbia, in 1988, and received M.S. ('91) and Ph.D. ('93) degrees in electrical engineering from Northeastern University, Boston, Massachusetts. She was a principal scientist at MIT, and in 2008 joined Northeastern University, where she is currently a professor of electrical and computer engineering. She is also a guest investigator at the Woods Hole Oceanographic Institution and a visiting scientist at MIT. Dr. Stojanovic is the recipient of the 2015 IEEE/OES Distinguished Technical Achievement Award. Her research interests include digital communications theory, statistical signal processing and wireless networks, and their applications to underwater acoustic systems. She is an Associate Editor for the *IEEE Journal of Oceanic Engineering* and a past Associate Editor for *IEEE Transactions on Signal Processing* and *IEEE Transactions on Vehicular Technology*. She also serves on the Advisory Board of *IEEE Communication Letters*, and chairs the IEEE Ocean Engineering Society's Technical Committee for Underwater Communication, Navigation and Positioning.

GAURAV S. SUKHATME [SM'05, F'11] is a professor of computer science at USC. He received his undergraduate education at IIT Bombay in computer science and engineering, and M.S. and Ph.D. degrees in computer science from USC. He is co-director of the USC Robotics Research Laboratory and director of the USC Robotic Embedded Systems Laboratory. His research interests are in robot networks with applications to environmental monitoring. He has served as Principal Investigator (PI) on numerous NSF, DARPA, and NASA grants. He is a co-PI on the Center for Embedded Networked Sensing, an NSF Science and Technology Center. He is a recipient of the NSF CAREER award and the Okawa Foundation research award. He is one of the founders of the Robotics: Science and Systems conference. He was Program Chair of the 2008 IEEE International Conference on Robotics and Automation and the 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems. He is Editor-in-Chief of *Autonomous Robots* and has served as Associate Editor of *IEEE Transactions on Robotics and Automation* and *IEEE Transactions on Mobile Computing*, and is on the Editorial Board of *IEEE Pervasive Computing*.